



Institute for Scientific Computing Research

# Seminar Series Abstracts

(in reverse chronological order)





September 29, 2000

# Multigrid Methods for Porous Media Flows

**Gabriel Wittum**

University of Heidelberg

Email: [wittum@iwr.uni-heidelberg.de](mailto:wittum@iwr.uni-heidelberg.de)

## *Abstract:*

**P**orous media flows occur in many situations, e.g., flow of ground and soil water, catalyzers, filters, flow through human tissue and skin, and many more. Modeling single and multi-phase flow through porous media is a challenging mathematical topic on account of their multi-scale character.

The simulation of such problems has to deal with many difficulties, such as heterogeneous materials resulting in stochastically jumping coefficients, non-linear problems with singular coefficient functions, complicated geometries, and, in particular for geological problems, strong anisotropies, convection dominated problems, and many more. All these problems are serious difficulties for multigrid methods. The most interesting topic in applying multigrid to porous media flow is how to relate the coarse-grid problems with the modelling scale.

In the lecture we show multigrid strategies for porous media flow problems. We show their successful application to a number of different applications.

*Research web page:* <http://www.iwr.uni-heidelberg.de/~techsim/>

*Institution web page:* [http://www.uni-heidelberg.de/index\\_e.html](http://www.uni-heidelberg.de/index_e.html)

September 28, 2000

*Abstract:*

# Algebraic Multigrid In An Industrial Environment

**Klaus Stueben**

GMD - Forschungszentrum  
Informationstechnik GmbH

Email: [stueben@gmd.de](mailto:stueben@gmd.de)

**A**lgebraic multigrid (AMG) has been shown to be very efficient and robust for the solution of various types of linear algebraic systems of equations, in particular those arising from the discretization of scalar partial differential equations. However, major research is still required to extend AMG's range of applicability. For instance, its applicability to systems of partial differential equations has not yet reached a robustness and efficiency comparable to that of the scalar case. But even for certain scalar applications, the performance of AMG may substantially deteriorate.

In principle, AMG is a highly interesting candidate for a "plug-in" solver in industrial codes. However, much work has to be invested to turn an academic code into one which satisfies industrial requirements. Rapid convergence is only one of these requirements. Much more important, however, are robustness and low memory requirements. The original AMG (AMG1R5) was far from really meeting these requirements to a satisfactory extent.

We will give a review on some of our experiences which led to the development of a completely new AMG code, SAMG. The underlying algorithm, although still close to the original AMG ideas, contains various enhancements and has been integrated as a solver into several commercial codes (from computational fluid dynamics, oil-reservoir simulation, and process simulation in semi-conductor applications). Developments regarding SAMG's integration into a commercial structural mechanics code are ongoing. We will discuss critical situations and present remedies for some particular scalar and systems applications.

*Speaker's web page:* <http://www.gmd.de/SCAI/people/stueben.html>

*Institution web page:* <http://www.gmd.de/Welcome.en.html>

September 27, 2000

*Abstract:*

# An Implicit, Nonlinear Newton–Krylov Resistive MHD Solver

**Luis Chacon**

Los Alamos National Laboratory

Email: [chacon@lanl.gov](mailto:chacon@lanl.gov)

Implicit time differencing of the resistive magnetohydrodynamic (MHD) equations can step over the limiting time scales—such as Alfvén time scales—to resolve the dynamic time scales of interest. However, nonlinearities present in these equations make an implicit implementation cumbersome. Here, viable paths for an implicit, nonlinear time integration of the MHD equations are explored using a 2D reduced viscous-resistive MHD model. The implicit time integration is performed using the Newton–Raphson iterative algorithm, employing Krylov iterative techniques for the required algebraic matrix inversions, implemented Jacobian-free (i.e., without ever forming and storing the Jacobian matrix). Convergence in Krylov techniques is accelerated by preconditioning the initial problem. A “physics based” preconditioner, based on an operator-split approximation to the original set of partial differential equations, is employed. The preconditioner employs low-complexity multigrid techniques to invert approximately the resulting elliptic algebraic systems. The resulting 2D reduced resistive MHD implicit algorithm is shown to be successful in dealing with large time steps (100–250 times the explicit Alfvén CFL limit) and fine grids (up to  $256 \times 256$ ). The algorithm is second-order accurate in time and efficient. Specifically, the number of Krylov iterations per Newton iteration scales very weakly with the total number of mesh points, and the CPU time scales as the time step to the  $-0.7$  power.

*Institution web page:* <http://www.lanl.gov/>



September 22, 2000

*Abstract:*

# Visualizing Multidimensional Geometry with Applications to Multivariate Problems

**Alfred Inselberg**

San Diego Supercomputer Center

Email: [aiisreal@math.tau.ac.il](mailto:aiisreal@math.tau.ac.il)

People working on multivariate (multidimensional) problems will benefit by understanding the underlying geometry, that is, learning what is possible and what is not. For example, in 1917 the physicist Ehrenfest showed that planetary orbits are stable only in dimension 3. Another dimensionality result is that rotating bodies have an axis of rotation only in odd-integer dimensions. The applications presented here will be more down to earth!

With a system of parallel coordinates a one-to-one mapping between subsets of  $N$ -space and subsets of 2-space is obtained. Lines in  $N$ -space are represented by  $N-1$  indexed points. In fact all  $p$ -flats (planes of dimension  $p$  in  $N$ -space) are represented by indexed points where the number of indices is one more than the object's dimension. The representations are generalized to enable the visualization of polytopes and certain kinds of hypersurfaces as well as Convex Objects in  $R^N$ . Synthetic constructions algorithms involving intersections, proximity, an interior point algorithm and "Line Topologies" of interest in Computer Vision will be presented. There will be interactive demonstrations of Biomedical applications, Process Control, Visual Data Mining (Yields in VLSI production, Finance, Retailing, Feature Extraction from LandSat Data etc.), then Collision Avoidance Algorithms for Air Traffic Control. A visual and computational model of a hypersurface suitable for Optimization, Trade-Off Analysis, and Decision support will be included.

*Speaker's web page:* [http://www.sdsc.edu/~nerona/Inselberg\\_A2/](http://www.sdsc.edu/~nerona/Inselberg_A2/)

*Institution web page:* <http://www.sdsc.edu/>

September 15, 2000

# Using the Forte Developer Performance Tools

**Marty Itzkowitz**

Sun Microsystems, Inc.

Email: [martyi@eng.sun.com](mailto:martyi@eng.sun.com)

## Abstract:

**T**he Forte Developer 6 performance tools are designed to help answer questions about application performance. In this talk we will discuss the kinds of performance questions that users typically ask, and the user-model for how the tools can help answer them. We will then discuss some specific type of questions for threaded applications, and for Parallel FORTRAN applications, and go through some examples of performance problems that the tools can help address.

For most applications, the basic question users have about performance is "What can I change to improve performance?" That translates into knowing what resources are being used, where in the program they are being used, and how the program got to that point. The design goal of the Forte Developer tools was to minimize the number of mouse-clicks to get the the "A-ha" point.

The user model for the tools is to build the application exactly as one would for production, including any optimization and parallelization. To see annotated source displays, use the `-g` flag for compilation (`-g` no longer disables optimization and parallelization). Then `dbx`, or the debugger GUI, can be used to run the application and measure its performance. From the recorded data the analyzer can present a list of functions, ordered by performance cost, and the callers and callees of any function, with their performance costs attributed to the callers and callees. The analyzer can also show annotated source displays, with per-line performance data, and annotated disassembly displays, with per-instruction data. In addition, the source and disassembly displays can show commentary from the compiler.

For multithreaded executables, users also want to know how the threads are being used, and what costs are involved in synchronizing them. For Parallel FORTRAN applications, the users also want to know which loops are parallel, and which serial, and what the obstacles are to parallelization of serial loops. The users are also interested in the efficiency of load-balancing in parallel regions. The tools provide data to answer all of these questions.

*Research web page:* <http://www.sun.com/forte/fortran/features.html>

*Institution web page:* <http://www.sun.com/>



September 13, 2000

*Abstract:*

# Direct Simulation of the High-Speed, Turbulent Reacting Shear Layer: Compressibility Effects

**Sutanu Sarkar**

Mechanical and Aerospace  
Engineering, UC San Diego

Email: [sarkar@mechanics.UCSD.edu](mailto:sarkar@mechanics.UCSD.edu)

In technological applications such as high-speed propulsion and energetic devices, the variation in thermodynamic variables associated with large heat release and/or high speeds interacts with and modifies the underlying turbulent flow. Direct numerical simulations of the reacting shear layer are performed over a wide range of heat release rates and convective Mach numbers to quantify and understand some of these modifications to the turbulence evolution and structure. Simulations are performed with up to 20 million grid points, high-order discretization, and with large evolution times required to achieve full-blown turbulence. The mixing rates and turbulence structure are found to be substantially different with respect to the incompressible case. DNS results and supporting analysis will be presented to explain the observed compressibility effects.

*Speaker's web page:*

[http://www.mae.ucsd.edu/GRAD\\_BROCH/fluids.htm#SARKAR/](http://www.mae.ucsd.edu/GRAD_BROCH/fluids.htm#SARKAR/)

*Institution web page:* <http://www.usd.edu/>

September 11, 2000

# Special Solution Strategies Inside a Spectral Element Ocean Model

**Gundolf Haase**

Johannes Kepler University, Linz

Email: [ghaase@numa.uni-linz.ac.at](mailto:ghaase@numa.uni-linz.ac.at)

## Abstract:

The weather forecast for a few days is rather accurate nowadays, but as soon as the prediction period reaches one month the accuracy of the prediction drops. This is because long-term weather is made in and above the ocean. Collecting more data from the atmosphere above the ocean would help to improve the forecast in coast regions, but long-term weather forecast also requires information on what's going on in the ocean.

I will focus my presentation on a ocean model that substitutes the 3D model by five 2D layers. Therein a filtering of the velocity vector requires the solution of the Laplacian equation in each layer for both velocity components. The discretization uses quadrilateral finite elements (cells) with 7th order test functions. The special difficulty therein consist in the huge amount of cells that restricts the range of applicable numerical methods with respect to memory usage.

I will present several techniques that could be used to solve the resulting system of equations by taking into account specialties of the cells. There are three approaches:

- Schur complement methods with a BPS-like preconditioner.
- Sparse approximation of the original matrix and use of algebraic multigrid methods.
- A Two- or Multigrid method with patch smoothing.

The element preconditioning technique by Stefan Reitzinger can be used in these approaches. Finally, the extension of the approaches to 3D is discussed.

*Speaker's web page:* <http://www.numa.uni-linz.ac.at/Staff/ghaase/haase.html>

*Institution web page:* <http://www.uni-linz.ac.at/>



August 28, 2000

# Cracking Multi-Representation Modeling

**Paul Reynolds**

University of Virginia

Email: [pfr@cs.virginia.edu](mailto:pfr@cs.virginia.edu)

## *Abstract:*

**P**lug-and-Play rarely works for multi-simulations—simulations composed of multiple stand-alone simulations. Among the obvious problems are semantically incompatible simulation components and mismatched units (e.g., of measure). We address a more difficult problem: maintaining a valid correspondence between representations of the same phenomenon in different component simulations when concurrent interactions occur. Correctness of a multi-simulation typically requires that a high degree of consistency be maintained between different representations.

Our approach has been to establish fundamental observations pertaining to the correctness of multi-simulations, and then to establish actions that support the construction of a correct multi-simulation, including the identification of mapping functions and the construction of attribute dependency graphs. We have defined a complete algorithm for constructing a correct multi-simulation. We have tested our approach on a small set of multi-simulations, mainly from the DoD, using each multi-simulation's respective DoD High Level Architecture Federated Object Model as our primary source material. Our approach has performed well in these tests.

We present our approach in depth, acknowledging that we have plucked mostly low-hanging fruit—generalizing and exploiting that which is more easily generalized and exploited. As a result, if you can't deliver mapping functions for maintaining consistency among overlapping representations, we can't necessarily deliver a correct multi-simulation. Are there properties of models that would allow a relaxation of our assumptions? Must every useful solution depend heavily on model semantics? The success of our presentation will be a function of the liveliness of the discussion it engenders on these and related questions.

*Speaker's web page:* <http://www.cs.virginia.edu/~pfr/>

*Research web page:* <http://www.cs.virginia.edu/~isotach/>

*Institution web page:* <http://www.virginia.edu/>

August 25, 2000

*Abstract:*

# Black Holes and Collaborative Computational Science

**Ed Seidel**

Max-Planck-Institut fuer  
Gravitationsphysik

Email: [eseidel@aei-potsdam.mpg.de](mailto:eseidel@aei-potsdam.mpg.de)

**L**arge-scale simulation problems, such as solving Einstein's equations of general relativity to simulate collisions of black holes, are driving the development of a new generation of advanced computational tools to enable scientists and engineers to work together to exploit "Grid-based" networks of parallel computers to solve problems on scales never before possible. I will discuss recent developments in this area, including the Cactus Computational Toolkit for parallel computing, the Globus Metacomputing Toolkit, and their applications to some important problems in astrophysics and relativity, including large-scale simulations of colliding black holes. Movies of these simulations will be shown. The techniques and toolkits discussed here are designed to address a large class of scientific and engineering applications.

*Research web page:* <http://www.cactuscode.org/>

*Institution web page:* <http://www.aei-potsdam.mpg.de/>



August 10, 2000

*Abstract:*

# Multi-Resolution Dynamic Meshes with Arbitrary Deformations

**Ariel Shamir**

University of Texas, Austin

Email: [arik@ticam.utexas.edu](mailto:arik@ticam.utexas.edu)

**M**ulti-resolution techniques and models have been shown to be effective for the display and transmission of large static geometric objects. Dynamic environments with internally deforming models and scientific simulations using dynamic meshes pose greater challenges in terms of time and space, and the need the development of similar solutions. In this talk I will introduce the T-DAG, a new adaptive multi-resolution representation for dynamic meshes with arbitrary deformations which include attribute, position, connectivity, and topology changes. T-DAG stands for Time-dependent Directed Acyclic Graph, which defines the structure supporting this representation. I will present an incremental algorithm (in time) for constructing the T-DAG representation of a given input mesh. This enables the traversal and use of the multi-resolution dynamic model for partial playback while still constructing new time-steps. The research is a joint work with V. Pascucci and C.L. Bajaj to appear on IEEE Visualization 2000.

*Speaker's web page:* <http://www.ticam.utexas.edu/~arik/>

*Institution web page:* <http://www.utexas.edu/>

August 9, 2000

*Abstract:*

# A Finite Difference Domain Decomposition Method Using Local Corrections for the Solution of Poisson's Equation

**Gregory Balls**

University of California, Berkeley

Email: [gballs@cs.berkeley.edu](mailto:gballs@cs.berkeley.edu)

A domain decomposition method will be presented for computing finite difference solutions to the Poisson equation with infinite domain boundary conditions. This method is a finite difference analogue of Anderson's Method of Local Corrections. The solution is computed in three steps. First, fine grid solutions are computed in parallel using infinite domain boundary conditions on each subdomain. Second, information is transferred globally through a coarse grid representation of the charge, and a global coarse grid solution is found. Third, a fine grid solution is computed on each subdomain using boundary conditions set with the global coarse solution, corrected locally with fine grid information from nearby subdomains.

There are three important features of this algorithm. First, the method requires only a single iteration between the local fine grid solutions and the global coarse representation. Second, the error introduced by the domain decomposition is small relative to the truncation error of a standard single-grid solution. Third, the computed solution is second-order accurate and only weakly dependent on the coarse grid spacing and the number of subdomains. As a result of these features, we are able to compute accurate solutions in parallel with a much smaller ratio of communication to computation than more traditional domain decomposition methods.

Results from implementations of this algorithm on an IBM SP2 and a Cray T3E will be presented. These results verify the overall accuracy and demonstrate the parallel scalability of the method.

*Speaker's web page:* <http://www.cs.berkeley.edu/~gballs/>

*Institution web page:* <http://www.berkeley.edu/>



August 8, 2000

*Abstract:*

# An Algebraic Multilevel Multigraph Algorithm

**Randy Bank**

University of California, San Diego

Email: reb@sdna1.ucsd.edu

**W**e describe an algebraic multilevel multigraph algorithm. Many of the multilevel components are generalizations of algorithms originally applied to general sparse Gaussian elimination. Indeed, general sparse Gaussian elimination with minimum degree ordering is a limiting case of our algorithm. Our goal is to develop a procedure which has the robustness and simplicity of use of sparse direct methods, yet offers the opportunity to obtain the optimal or near-optimal complexity typical of classical multigrid methods.

*Speaker's web page:* <http://scicomp.ucsd.edu/~reb/>

*Institution web page:* <http://www.ucsd.edu/>

August 1, 2000

*Abstract:*

# The Constraints in the Einstein Equations: Well-posedness, a Priori and a Posteriori Error Estimates, and Some Numerical Solutions

**Michael Holst**

University of California, San Diego

Email: mholst@math.ucsd.edu

In this talk we consider a coupled nonlinear elliptic system representing the Hamiltonian and momentum constraints in the Einstein equations. This system must be solved exactly or numerically to produce consistent initial data for general relativistic simulations of black hole and neutron star collisions. Moreover, the constraints must hold at all times in dynamical situations. Well-posedness of the system on connected compact Riemannian manifolds with Lipschitz boundaries was previously unstudied, and therefore in the first part of the talk we establish that the constraints have unique weak solutions under minimal smoothness assumptions on the data. The proof technique (Riesz-Schauder theory and convex analysis) allows for some degree of negative conformal scalar curvature. We also establish two quasi-optimal a priori error estimates for Galerkin approximations, and we derive an a posteriori error estimate which leads to two distinct error indicators for adaptive simplex subdivision algorithms (one estimator is purely residual based, the other involves a linearized dual problem).

In the second part of the talk we outline an implementation of the adaptivity techniques using the adaptive finite element software package "MC," which is designed to adaptively solve general nonlinear systems of tensor equations on manifolds. MC is an adaptive multilevel finite element package which employs a posteriori error estimation, simplex subdivision, algebraic multilevel methods, global inexact Newton methods, and numerical continuation methods, for the solution of coupled elliptic systems on 2- and 3-manifolds. We demonstrate the capabilities of MC by generating initial data for several different types of black hole spacetimes, including a two-hole collision. To perform these calculations, MC uses a completely new low-communication algorithm for parallel adaptive finite element methods, which was developed jointly with R. Bank at UCSD.

*Speaker's web page:* <http://scicomp.ucsd.edu/~mholst/>

*Institution web page:* <http://www.ucsd.edu/>



July 25, 2000

*Abstract:*

# Modelling Axisymmetric Laminar Flames: An Application of Local Rectangular Refinement Solution-Adaptive Gridding

**Beth Anne Bennett**

Yale University

Email: [beth.bennett@yale.edu](mailto:beth.bennett@yale.edu)

Over the past two decades, scientific interest in multidimensional laminar and turbulent combustion phenomena has greatly increased, mainly because of tightened government restrictions on pollutant emissions, along with a growing public awareness of environmental concerns. Although today's computers are still not powerful enough to solve exact numerical models of three-dimensional, unsteady, turbulent combustion, a class of problems which *can* be solved is that of chemically complex, multidimensional laminar flames. Because reaction zone thicknesses in such flames can be less than one-hundredth of the overall flame height, the demands on CPU time and memory can be drastically reduced (and the computations made feasible) with the application of a sophisticated adaptive gridding technique.

The results presented here have been obtained using the local rectangular refinement (LRR) solution-adaptive gridding method, which automatically produces orthogonal unstructured adaptive grids and employs multiple-scale finite differences to discretize the elliptic governing PDEs. These strongly coupled, highly nonlinear discretized equations are then solved simultaneously via Newton's method with a nested Bi-CGSTAB solver. The unstructured nature of the grids has necessitated substantial modification of the standard solver. Comparisons with computed results on equivalently refined conventional grids indicate that the LRR method provides substantial savings in execution time and computer memory requirements, with little or no loss in solution accuracy.

Three types of steady-state axisymmetric laminar methane/air flames are examined: a diffusion flame; a set of six partially premixed flames of various primary equivalence ratios; and a pair of lean and rich Bunsen flames. Each flame model employs a vorticity-velocity formulation of the governing equations and includes GRI-Mech chemistry, multicomponent transport, and optically thin radiation modelling. Agreement with available experimental measurements of temperature and species concentrations ranges from very good to excellent.

*Speaker's web page:* <http://www.eng.yale.edu/faculty/vita/bennettbeth.html>

*Institution web page:* <http://www.yale.edu/>

July 24, 2000

*Abstract:*

# First-Order System $LL^*$ (FOSLL $^*$ ): Scalar Elliptic Partial Differential Equations

**Zhiqiang Cai**

Purdue University

Email: [zcai@math.purdue.edu](mailto:zcai@math.purdue.edu)

The  $L^2$ -norm version of first-order system least squares (FOSLS) attempts to reformulate a given system of partial differential equations so that applying a least-squares principle yields a functional whose bilinear part is  $H^1$ -elliptic. This means that the minimization process amounts to solving a loosely coupled system of elliptic scalar equations. An unfortunate limitation of the  $L^2$ -norm FOSLS approach is that this product  $H^1$  equivalence generally requires sufficient smoothness of the original problem. Inverse-norm FOSLS overcomes this limitation, but at a substantial loss of real efficiency. The FOSLL $^*$  approach introduced here is a promising alternative that is based on recasting the original problem as a minimization principle involving the adjoint equations. This paper provides a theoretical foundation for the FOSLL $^*$  methodology and illustrates its performance by applying it numerically to several different examples. Results for the so-called two-stage approach applied to discontinuous coefficient problems show efficiency that is much better than inverse-norm least-squares methods. FOSLL $^*$  appears to exhibit the generality of the inverse-norm FOSLS approach while retaining the full efficiency of the  $L^2$ -norm approach. This is a joint work with Manteuffel, McCormick, and Ruge at University of Colorado at Boulder.

*Speaker's web page:* <http://www.math.purdue.edu/~zcai/>

*Institution web page:* <http://www.purdue.edu/>



July 19, 2000

*Abstract:*

# Multigrid for 2-D High Reynolds Number Turbulent Transonic Flows

**Steven Allmaras**

Boeing

Email: [steven.r.allmaras@boeing.com](mailto:steven.r.allmaras@boeing.com)

**W**e present multigrid methods for the solution of the steady-state 2-D compressible Navier–Stokes equations discretized by an upwind finite-volume scheme on structured grids. A one-equation turbulence model is included in the set of equations. We describe the appropriate multigrid ingredients necessary to achieve effective convergence rates for high Reynolds number turbulent transonic flows. We describe trades for the formulation of the equations to be relaxed, the multigrid coarsening strategy, the choice of relaxation on each grid, and the degree of coupling between the conservation equations and the turbulence model equation. We also propose modifications for the one-equation turbulence model to alter its non-physical transient behavior as an aid to achieve steady-state solutions; these modifications, however, do not change the steady-state formulation of the model.

We present results for two cases: an inviscid transonic channel flow and a turbulent transonic airfoil. The achieved convergence rates are substantial improvements over those obtained by state-of-the-art methods.

We finish with a discussion of the extension of the method to 3-D structured grids and possibilities for extension to unstructured grids. Issues related to parallel implementation are also discussed.

*Institution web page:* <http://www.boeing.com/>

June 27, 2000

*Abstract:*

# Liquid Crystalline Polymers and Other Macromolecular Fluids

**Gregory Forest**

University of North Carolina

Email: [forest@amath.unc.edu](mailto:forest@amath.unc.edu)

Many high-performance materials are made from macromolecular fluids, ranging from spider silk to synthetic liquid crystalline polymers like Kevlar and Vectra. Flows of these materials couple the dynamics of anisotropic macromolecules with hydrodynamics. There are many length-scales, and approaches, in the modeling of these systems. We have focused on a mesoscale description following Doi & Edwards, in terms of an orientation tensor for the second moment of the molecules with respect to a probability distribution function. This object relates directly to measurements of quantities such as birefringence or light scattering intensities. We will describe various phenomena we have studied and advertise some interesting challenges. Topics include a revisitation of the classical isotropic-to-nematic phase transition, what happens to this transition in the presence of imposed flows, and methods to construct heterogeneous patterns routinely observed in experiments. Collaborators include Qi Wang, IUPUI, and Hong Zhou, UC-Santa Cruz.

*Speaker's web page:* <http://www.amath.unc.edu/Faculty/forest/>

*Research web page:* <http://www.amath.unc.edu/>

*Institution web page:* <http://www.unc.edu/>

June 14, 2000

*Abstract:*

# On the Analysis and Construction of Absorbing Layers in CEM and CFD

**Saul Abarbanel**

Tel Aviv University

Email: saula@math.lbl.gov

**T**his talk will review some recent results about the theory and application of PML (Perfectly Matched Layers):

1. It will be shown that that set of split PDE's of the Berenger PML method is not strongly well-posed. Under certain perturbations its solutions may be inappropriate.
2. A recently introduced system of PDE's (Zhao & Cangellaris, 1996; Ziolkowski, 1997; Petropoulos et al., 1998), based on physical considerations, and which describes the behavior of EM waves in artificial absorbing layers, is analyzed. Analytic solutions are found for the cases of semi-infinite and finite depth layers, both for primitive and characteristic boundary conditions.

A different set of equations, based on mathematical considerations, and which offers some advantages, is constructed and analyzed for the same geometries and boundary conditions.

Both sets of PDE's are strongly well posed.

3. A set of strongly well-posed PML equations is developed for describing the absorption of acoustic and vorticity waves in two-dimensional convective aeroacoustics under the assumption of a constant mean flow. A central piece in the derivation is the development of a coordinate transformation that conserves the dispersion relation from the static acoustic case. Numerical solutions illustrate the efficacy of this PML scheme even when applied to the fully nonlinear Euler equations.

This work was done in collaboration with David Gottlieb and Jan Hesthaven of Brown University.

*Institution web page: <http://www.tau.ac.il/>*

June 12, 2000

# Using a Multilayer Ocean Model on Clusters versus a Traditional Supercomputer

**Craig Douglas**

University of Kentucky

Email: [douglas@ccs.uky.edu](mailto:douglas@ccs.uky.edu)

## *Abstract:*

**W**e simulate oceanic overflow problems. The long-term goal of this research is to understand both the local dynamics of downslope flows in the ocean and their role in the Earth's global thermohaline circulation. The modeling of these flows and their climatic impact is complicated by the inherent range of spatial scales involved, which extend from the global scale [O(10,000) km] down to the local scale of the overflows themselves [O(1) km], and by the intrinsic three dimensionality of the overflow dynamics.

The Spectral Element Ocean Model (SEOM) offers an elegant solution to these difficulties. It features advanced algorithms, based on  $h$ - $p$  type finite element methods, allowing accurate representation of complex coastline and oceanic bathymetry, variable lateral resolution, and high-order solution of the three dimensional oceanic equations of motion.

SEOM's geometrical flexibility permits highly inhomogeneous horizontal grids. An added advantage of the technique is its scalability. Most of the computations are carried out at the element level; only interface information needs to be exchanged between elements. The dual characteristic of dense and structured local computations, and sparse and unstructured communication, enhances the locality of the computations, and makes SEOM ideally suited for parallel computers.

In this talk we demonstrate what types of cluster machine characteristics are suitable for solving our problems, how much they cost at what point in time, and how they compare to a traditional RISC based supercomputer solution. Our comparisons will note time, money, and aggravation factors.

*Speaker's web page:* <http://www.ccs.uky.edu/~douglas/>

*Research web page:* <http://www.mgnet.org/>

*Institution web page:* <http://www.uky.edu/>



June 12, 2000

# Random Testing with "Fuzz": A Decade of Random Testing

**Barton Miller**

University of Wisconsin

Email: [bart@cs.wisc.edu](mailto:bart@cs.wisc.edu)

## *Abstract:*

In 1990, we published the results of a study of the reliability of standard UNIX application/utility programs. This study showed that by using simple (almost simplistic) random testing techniques, we could crash or hang 25%–33% of these utility programs. In 1995, we repeated and significantly extended this study using the same basic techniques: subjecting programs to random input streams. This study also included X-Window applications and servers. A distressingly large number of UNIX applications still crashed with our tests. The X-window applications were at least as unreliable as command-line applications. The commercial versions of UNIX fared slightly better than in 1990, but the biggest surprise was that Linux and GNU applications were significantly more reliable than the commercial versions.

This year (2000), we took another stab at random testing, this time testing applications running on Windows/NT. Given valid random mouse and keyboard input streams, we could crash or hang 45% of these applications.

In this talk, I will discuss our simple testing procedure as it applies to UNIX and Windows. I will then present the UNIX and NT test results. These results include, in many cases, identification of the bugs and the coding practices that caused the bugs. In several cases, these bugs introduce issues relating to system security. The talk will conclude with some philosophical musings on the current state of commercial software.

Papers on the three studies (1990, 1995, and 2000) can be found at:  
[ftp://grilled.cs.wisc.edu/technical\\_papers/fuzz.ps](ftp://grilled.cs.wisc.edu/technical_papers/fuzz.ps)  
[ftp://grilled.cs.wisc.edu/technical\\_papers/fuzz-revisited.ps](ftp://grilled.cs.wisc.edu/technical_papers/fuzz-revisited.ps)  
[ftp://grilled.cs.wisc.edu/technical\\_papers/fuzz-nt.ps](ftp://grilled.cs.wisc.edu/technical_papers/fuzz-nt.ps)

PDF versions are also available (change the .ps to .pdf). The tools used for the testing can be found at:  
<ftp://grilled.cs.wisc.edu/fuzz/>

*Institution web page:* <http://www.wisc.edu/>

June 8, 2000

# Sputnik: A Programming Model for Automated Decomposition on Heterogeneous

**Sean Peisert**

San Diego Supercomputer Center

Email: [peisert@sdsc.edu](mailto:peisert@sdsc.edu)

## *Abstract:*

Clusters of multiprocessor nodes are becoming common in scientific computing. As a result of the expandability of clusters, faster nodes are frequently added and older nodes are gradually removed, making the cluster heterogeneous. As heterogeneity increases, traditional methods for programming clusters of multiprocessors become less optimal, because they do not account for the fact that a cluster will only run as fast as the slowest node. Sputnik is a programming methodology and software library that addresses the problem of heterogeneity on a dedicated cluster of multiprocessors.

Sputnik uses a two-stage process for running applications on a cluster of multiprocessors. The first stage assesses the relative performance of each node by running the program individually on each node, determining from the run times both the performance and application-specific optimization. Using the timings obtained from stage one, the second stage partitions the dataset non-uniformly, according to the relative speed of each node. All future runs of the program use the optimal partitionings and number of threads per node.

Sputnik is implemented on top of the KeLP infrastructure to handle irregular decomposition and data motion. It enables code to be written for a heterogeneous cluster as if the cluster is homogeneous. Sputnik can run scientific applications on a heterogeneous cluster faster, with improved utilization, than a nearly identical program written in KeLP alone. Experimental results from a pair of SGI Origin2000's indicate that Sputnik can improve run-time of an iterative solver for Poisson's equation by 35%.

*Speaker's web page:* <http://www.sdsc.edu/~peisert/>

*Institution web page:* <http://www.sdsc.edu/>



June 5, 2000

*Abstract:*

# Development and Application of a Parallel, Solution- Adaptive Code for Space Physics

**Kenneth Powell**

University of Michigan

Over the past several years, a multidisciplinary group in space physics, CFD and computer science has worked to develop a solution-adaptive, parallel code for space physics. The uses dynamic adaptation to resolve the multiple scales inherent in this class of problems, and has achieved high parallel efficiency on a number of machines. On a 1490-processor Cray T3E, speeds above 300 GigaFLOPs have been sustained, even while using the dynamic adaptation abilities of the code. The development of the current code, some applications of the current code, and plans for future work will be discussed.

*Speaker's web page:* <http://www-personal.engin.umich.edu/~powell/>

*Research web page:*  
<http://www.engin.umich.edu/dept/aero/people/faculty/powell.html>

*Institution web page:* <http://www.umich.edu/>

May 31, 2000

*Abstract:*

# Large Multiscale Tensor-Product Finite Element Approximation Spaces

**Stephan Knapek**

University of Bonn

Email: knapek1@lnl.gov

**D**espite the rapid development of computer technology, many problems in scientific computation remain intractable to a direct numerical simulation due to storage requirements and numerical complexity. Especially problems in higher dimension remain well out of the range of today's largest computers even after the use of parallelization. This intractability of higher-dimensional problems is often referred to as the "Curse of dimension."

The situation looks different when we impose specific regularity assumptions on our problem. Then, Finite Element subspaces with relatively large dimensions that contribute only "little" to the error reduction can be identified and omitted from the approximation space. This idea has appeared under various names (boolean blending schemes, hyperbolic cross points, sparse grids, ...) in approximation and interpolation theory and attracted much attention in recent time.

In this talk we generalize the concept of hyperbolic cross points and discuss applications to operator equations such as integral and partial differential equations. We construct a scale of nested finite element spaces based on multiscale tensor-product bases, and we show under which circumstances these approximation spaces break the curse of dimension.

*Institution web page:* [http://www.uni-bonn.de/unibonn\\_en.html](http://www.uni-bonn.de/unibonn_en.html)



May 25, 2000

# MPI Software Technology's Path Forward Program for Terascale MPI Deployment on ASCI Architectures

**Anthony Skjellum**

MPI Software Technologies Inc.

Email: [tony@mpi-softtech.com](mailto:tony@mpi-softtech.com)

## *Abstract:*

In this presentation, we discuss the technology underpinnings and plans for achieving a new-generation of ASCI-relevant MPI implementations, based on the Software Pathforward Program. MPI Software Technology, Inc. (MSTI) is providing improved MPI-1 and MPI-2 capabilities for tri-lab ASCI systems, based on its existing and developing commercial technology for high performance messaging middleware. Specific technological achievements, including thread safety, high performance derived data types, and overlapping of communication and computation, are mentioned. This talk also provides background on the company, its middleware products/technologies, and work in clusters and embedded MPI deployments. Strategies for moving MPI/Pro—the commercial-grade MPI offered by MSTI—to ASCI platforms is included.

At the conclusion, input from users about deficiencies in existing ASCI and cluster MPI settings is invited, as is input about MPI-2 perceptions/interest/priorities moving forward.

*Speaker's web page:* <http://www.mpi-softtech.com/company/president.html>

*Research web page:* <http://www.mpi-softtech.com/>

May 18, 2000

# High Resolution Central Schemes

**Eitan Tadmor**

University of California, Los Angeles

Email: [tadmor@math.ucla.edu](mailto:tadmor@math.ucla.edu)

## Abstract:

**W**e discuss recent developments of high-resolution schemes for the approximate solution of hyperbolic conservation laws, Hamilton–Jacobi (HJ) equations and related nonlinear problems.

We focus on non-oscillatory central schemes as prototype for Godunov-type projection methods. A variety of numerical experiments demonstrate that the proposed central schemes offer simple, robust, Riemann-solver-free “black box” solvers, while at the same time they retain the high-resolution content of the more expensive upwind schemes.

Among several issues to be discussed will be the following highlighted features:

**Scalar equations.** Variation and entropy stability estimates as well as multidimensional L1-bounds are presented. New convergence results based on both global and pointwise error estimates follow.

**Systems of equations.** Extension to systems is carried out by componentwise application of the scalar framework. It is in this context that the central schemes offer a remarkable advantage over the corresponding upwind framework.

**Multidimensional problems.** Since we bypass the need for (approximate) Riemann solvers, multidimensional problems are solved without dimensional splitting. In fact, the proposed class of central schemes is utilized for a variety of nonlinear transport equations, and in this context we demonstrate the construction and implementation of central schemes for HJ and incompressible Euler equations.

We describe recent developments of new high-resolution central schemes for general convection-diffusion equations.

We overview recent applications to various models, including incompressible flows, MHD equations, simulations of semi-conductors models, and more.

*Speaker's web page:* <http://www.math.ucla.edu/~tadmor/>

*Research web page:* <http://www.math.ucla.edu/faculty/tadmor.html>

*Institution web page:* <http://www.ucla.edu/>

May 18, 2000

*Abstract:*

# UCLA's new NSF Institute for Pure and Applied Mathematics: Program and Collaborations with LLNL

**Eitan Tadmor**

University of California, Los Angeles

Email: [tadmor@math.ucla.edu](mailto:tadmor@math.ucla.edu)

UCLA will host one of NSF's three mathematics institutes as of September 2000. The first-year program will consist of units on Functional Genomics, Financial Mathematics, Oscillatory Integrals and Dispersive Equations, and Geometrically Based Motions. An overview of the institute and of opportunities for LLNL-IPAM collaborations will be presented.

*Speaker's web page:* <http://www.math.ucla.edu/~tadmor/>

*Research web page:* <http://www.math.ucla.edu/faculty/tadmor.html>

*Institution web page:* <http://www.ucla.edu/>

May 9, 2000

# Sparse Pattern Selection Strategies for Robust Frobenius Norm Minimization

**Iain Duff**

CCLRC Rutherford Appleton  
Laboratory

Email: [I.S.Duff@rl.ac.uk](mailto:I.S.Duff@rl.ac.uk)

## *Abstract:*

**W**e consider preconditioning strategies for the iterative solution of dense complex symmetric non-Hermitian systems arising in computational electromagnetics. We consider in particular sparse approximate inverse preconditioners that use a static nonzero pattern selection. The novelty of our approach comes from using a different nonzero pattern selection for the original matrix than for the preconditioner and from exploiting geometrical information from the underlying meshes in addition to using methods based on the magnitude of the entries.

The numerical and computational efficiency of the proposed preconditioners are illustrated on a set of model problems arising both from academic and from industrial applications.

The results of our numerical experiments suggest that the new strategies are viable approaches for the solution of large-scale electromagnetic problems using preconditioned Krylov methods. In particular, our strategies are applicable when fast multipole techniques are used for the matrix-vector product on parallel distributed memory computers.

*Speaker's web page:* <http://www.cse.clrc.ac.uk/Person/I.S.Duff/>

*Institution web page:* <http://www.rl.ac.uk/>



April 28, 2000

*Abstract:*

# Numerical Simulation Methods for Electromagnetic Scattering Problems

**Karl Warnick**

University of Illinois, Urbana-  
Champaign

Email:  
warnick@sunchew.ece.uiuc.edu

**E**lectromagnetic analysis of radar targets, antennas, microwave circuits, digital systems, communications links, and optical devices is a significant challenge in the field of scientific computation. Due to the indefinite, wavelike character of these problems, efficient methods such as multigrid, which are highly successful for static and electrically small problems, break down as electrical size increases. For surface integral equation or boundary element solvers, the Fast Multipole Method (FMM) has brought dramatic progress towards attaining  $O(1)$  steps per degree of freedom. Three outstanding problems remain: grid robustness, accuracy control, and increasing computational efficiency in the high-frequency asymptotic limit. This presentation will explore these open questions and will outline current research directions, including methods for integrating multigrid with FMM in the high frequency regime.

*Research web page:* <http://www.ccem.uiuc.edu:80/reschew.html>

*Institution web page:* <http://www.uiuc.edu/>

April 25, 2000

# Scalable Adaptive Algorithms, Scalable Platforms

**Gerhard Zumbusch**

University of Bonn

Email: [zumbusch@iam.uni-bonn.de](mailto:zumbusch@iam.uni-bonn.de)

## *Abstract:*

**W**e present some highlights of the parallel computing activities of our research group. These include the construction of a parallel computing platform, namely our Parnass2 PC cluster as Europe's fastest cluster so far, and the development of several parallel multilevel and multigrid algorithms, especially in the presence of adaptive grid refinement, and the related parallel load-balancing issues.

The cluster Parnass2 at the Institute for Applied Mathematics at the University of Bonn was designed and built by the group of researchers. The 144-processor cluster is listed on both the current and the previous TOP500 list of the world's supercomputers. It is listed there among only two further clusters, namely the larger CPlant at Sandia (Albuquerque) and the Avalon cluster at LANL, both in the United States. Hence, the Parnass2 cluster is probably the cheapest computer on the lists. Together with the superior parallel performance of the Linpack benchmark (even among commercial clusters) it shows best price-performance ratio. We will discuss and compare several issues that lead to the current design, including network performance and software and SMP computing nodes. The comparison both to the Cray T3E and to IBM SP series (and to ASCI blue pacific) are especially of interest.

The second part of the talk is devoted to algorithmic developments. Here we concentrate on parallel implementations of adaptive grid refinement and the related load balancing issue. We propose space-filling curve partition methods for locality both on the message-passing level and within the memory of a single node. We discuss iterative multilevel and multigrid implementations on top of such grids with grid-independent convergence rates and we highlight how the partition method affects the complexity of the implementation and the performance of such a PDE solver. Furthermore, alternative discretization schemes based on adaptive wavelet methods and their impact on parallelization are discussed.

*Speaker's web page:* <http://wissrech.iam.uni-bonn.de/people/zumbusch.html>

*Research web page:* <http://wissrech.iam.uni-bonn.de/>

*Institution web page:* [http://www.uni-bonn.de/unibonn\\_en.htm](http://www.uni-bonn.de/unibonn_en.htm)



April 13, 2000

*Abstract:*

# Techniques Leading to Scalable Algorithms for Sparse Linear Systems

**X. Sherry Li**

Lawrence Berkeley Laboratory

Email: [xiaoye@nslsc.gov](mailto:xiaoye@nslsc.gov)

We present some useful techniques for sparse direct and iterative solvers on large-scale parallel machines. Many classical algorithms are re-designed and new algorithms developed to enhance scalability. One example is sparse LU factorization with partial pivoting. It is hard to implement efficiently on distributed memory machines, because of its somewhat unpredictable way of generating fine-grained work and intermediate results at runtime. We have investigated the possibility of replacing (dynamic) partial pivoting by some other techniques to control the element growth during the elimination. These include statically pre-pivoting large elements to the diagonal and dynamically modifying the small diagonal pivot. These numerical techniques enable us to achieve a more scalable MPI implementation of the LU factorization.

Another example is ordering of the equations and variables of a sparse linear system. For a direct solver, the goal of ordering is usually to minimize the fills in the factored matrices. But the best fill-reducing ordering may not give the best parallel runtime. For a CG-like solver, the ordering has great impact on the efficiency of the matrix-vector product kernel. We will illustrate these by comparing several popular ordering strategies.

*Speaker's web page:* <http://www.nslsc.gov/~xiaoye/>

*Research web page:* <http://www.nslsc.gov/>

*Institution web page:* <http://www.lbl.gov/>

April 6, 2000

# Probabilistic Clustering of Sequences, Trajectories, and other Non-Vector Data

**Padhraic Smyth**

University of California, Irvine

Email: [smyth@ics.uci.edu](mailto:smyth@ics.uci.edu)

## *Abstract:*

Clustering has a long history in exploratory data analysis. Searching for homogenous data clusters is a fundamental component of any endeavour to discover structure from data. This talk will begin with a brief overview of a general model-based probabilistic clustering framework. This approach, relatively well known in applied statistics, focuses on the use of mixture models as an underlying generative model for observed data. It also provides a relatively objective and sound methodology for determining the model which is closest to truth.

A particularly useful aspect of the probabilistic approach is the ability to generalize from clustering in vector-spaces to clustering sequences, trajectories, and other “dynamic” data that we commonly observe from individuals and/or systems. The main part of the talk will introduce a new and general framework for mixture-model clustering of such data. The probabilistic approach solves in an elegant and coherent manner the dual difficulties of (a) how to define distance metrics between observations (e.g., sequences) of different sizes, and (b) how to “weight” different individuals for whom we have different amounts of data. Illustrative applications include unsupervised learning of gestures from video data, clustering of individuals based on Web browsing behavior, modeling of gene expression data, modeling of cyclone trajectories, clustering locusts based on observed motor behavior, and clustering of medical patients based on histograms of red blood cells. A subset of these applications will be discussed, time permitting. The talk will conclude with a brief discussion of some apparently useful connections between mixture modeling in this context and Bayesian hierarchical models.

*Speaker's web page:* <http://www.ics.uci.edu/~smyth/>

*Institution web page:* <http://www.uci.edu/>

March 29, 2000

*Abstract:*

# Large-Scale PDE-Constrained Optimization: Parallel Algorithms and Applications to Optimal Design, Optimal Control, and Parameter Estimation Problems

**Omar Ghattas**

Carnegie Mellon University

Email: [oghattas@cs.cmu.edu](mailto:oghattas@cs.cmu.edu)

**V**ery large scale PDE-constrained optimization problems arise naturally in many areas of engineering and science, often taking the form of optimal design, optimal control, or parameter estimation problems. The common denominator is a nonlinear optimization problem that is constrained by the PDEs that govern behavior of the physical system. Thus, solving the PDEs is just a subproblem associated with optimization, which can be orders of magnitude more challenging computationally. Despite its importance, little attention has been devoted to the design of parallel algorithms for PDE-constrained optimization. This is expected: it makes little sense to address the “inverse” problem until the “forward” problem is well understood. However, recent advances in parallel PDE solvers and the arrival of the teraflop computing era motivate the pursuit of very large-scale simulation-based optimization. Despite these advances, many PDE-constrained problems remain intractable with current optimization technology. To render them tractable, we need to develop parallel optimization algorithms that exploit the PDE nature of the constraints, scale to the millions of constraints and variables that arise upon discretization, and capitalize on emerging highly parallel supercomputers.

I will give an overview of the TAOS (Terascale Algorithms for Optimization of Simulations) Project at CMU, whose goals are to develop the enabling parallel numerical algorithms for large-scale simulation-based optimization, and to apply them to driving optimal control, optimal design, and parameter estimation problems in engineering and science. The talk will be illustrated with motivating applications in each of the three classes: optimal design of artificial heart devices, optimal boundary control of viscous flows, and inverse earthquake ground motion modeling. I will describe a family of parallel Newton–Krylov methods for solution of the optimality system of PDE-constrained optimization problems. These are essentially full-space Newton-Sequential Quadratic Programming (SQP) methods with preconditioning by reduced-space limited memory quasi-Newton SQP and approximate forward Jacobians. This combines the rapid convergence of Newton methods with the low per-iteration cost of approximate methods. I will present studies of parallel efficiency and scalability of a PETSc-based implementation for the problem of optimal control of a viscous incompressible fluid by suction/injection of fluid on its boundary. Numerical experiments for problems of size up to 1.5 million state variables and 50,000 control variables on up to 256 processors of a Cray T3E-900 yield encouraging results. I will also discuss a space-and-time-efficient adjoint algorithm for time dependent PDE-constrained optimization, with applications to inverse wave propagation.

The TAOS Project is jointly directed with Larry Biegler. The research on Newton–Krylov optimization methods and optimal flow control is joint work with Ph.D. student George Biros; inverse earthquake modeling with Ph.D. students Volkan Akcelik and Yiannis Epanomeritakis and collaborator Jacobo Bielak; and artificial heart design with Ph.D. student Ivan Malcevic and collaborators Jim Antaki and Greg Burgreen (University of Pittsburgh Medical Center).

*Speaker's web page:* <http://www.cs.cmu.edu/~oghattas/omar.html>

*Institution web page:* <http://www.cmu.edu/>

March 28, 2000

# Stable Embedded Finite-Difference Time-Domain Methods For Maxwell Equation

Jan Hesthaven

Brown University

Email: Jan\_Hesthaven@brown.edu

## Abstract:

The last decade has seen a renewed interest in the development of flexible, simple, and accurate methods for solving Maxwell's equations. These efforts have partly been guided by technological developments requiring the ability to model broadband signals and their interaction with very complex configurations of materials. To enable such flexibility much effort has been centered around the development of time-domain methods in which Maxwell's equations are solved as a hyperbolic system.

The most popular time-domain method currently used in electromagnetic computations is based on central second-order finite differences in space and a Leapfrog scheme in time implemented on Cartesian grids, staggered in time and space, and the computational problem is simply embedded into this grid by assigning different material properties to the grid points. The advantages of this approach, known as the Yee scheme in computational electromagnetics, are its simplicity and robustness which make it the favorite tool for practitioners. However, this approach suffers from a number of significant problems such as a need to employ staircase approximations to curved boundaries and interfaces and an inability to enforce the correct electromagnetic boundary conditions across material interfaces.

In this presentation, we discuss a stable second-order Cartesian grid finite difference scheme for the solution of Maxwell's equations. The scheme employs a staggered grid in space while several choices are possible for the temporal integration. However, contrary to the Yee scheme, this new scheme represents the physical location of the material and metallic boundaries correctly to the order of the scheme, hence eliminating the problems caused by staircasing, and enforces the correct jump-conditions on the individual field components across material interfaces.

Accuracy and strict stability are discussed for the one-dimensional scheme and issues related to two- and three-dimensional problems are briefly mentioned. The analysis exposes that the effects of staircasing as well as a lack of properly enforced jump-conditions on the field components have a significant impact on the global accuracy. Indeed, the analysis reveals that for cases where a field component is discontinuous along a gridline, as happens for general two- and three-dimensional curvilinear material interfaces, the Yee scheme may exhibit local divergence and loss of global convergence.

This new formulation, the computational cost of which is primarily incurred during a preprocessing stage, also lends itself to a very efficient parallel implementation by allowing for the definition of 'ghost'-interfaces to minimize the communication costs.

To validate the analysis several one- and two-dimensional test cases are presented, showing an improvement of typically 1-2 orders of accuracy at little or no additional computational cost over the Yee scheme, which in general remains first order accurate at best.

The work discussed is done in collaboration with Dr. Adi Ditkowski (Brown University), Dr. Kim Dridi (Ris National Laboratory, Denmark), and Mr. Chun-Hau Teng (Brown University).

*Speaker's web page:* <http://www.cfm.brown.edu/people/jansh/>

*Research web page:* <http://www.cfm.brown.edu/home.html>

*Institution web page:* <http://www.brown.edu/>

March 27, 2000

*Abstract:*

# High-Order Unstructured Grid Methods for Conservation Laws

**Jan Hesthaven**

Brown University

Email: [Jan\\_Hesthaven@brown.edu](mailto:Jan_Hesthaven@brown.edu)

**I**ncreasing interest in the modeling of unsteady problems over very long times and in computationally very large domains requires that accurate and efficient high-order methods be developed for such problems. The geometric complexity, however, of problems of scientific and industrial interest makes such developments a very significant challenge.

The application of high-order methods, and in particular spectral methods, is traditionally restricted to domains that can be smoothly mapped to a unit cube, hence allowing for the construction of a well-behaved multi-dimensional approximation using tensorproducts. Unfortunately, as is well known, automated grid generation and adaptive meshing is greatly complicated by using only hexahedrals as the fundamental element.

Guided by these observations we shall discuss how to develop accurate high-order/spectral methods on triangles and tetrahedra in which the approximation is based upon truly multi-dimensional Lagrange interpolation polynomials much in the spirit of classical spectral collocation methods. To enable such a construction we first discuss the construction of families of nodal sets well suited for high-order interpolation on the  $n$ -simplex and show how to address this issue through the solution of problems in generalized electrostatics. Among other things we emphasize how one can exploit the symmetries in the nodal sets to factorize the discrete operators into a sparse format, hence allowing for a fast evaluation of matrix-vector products to approximate spatial derivatives.

For the approximation of initial boundary value problems we shall discuss how to construct stable methods on almost general nodal distributions. The central problem is how to impose the boundary conditions in a stable way - an issue we resolve by showing how to imposing the boundary conditions only weakly through a penalty term. As we shall show, the resulting schemes contain, as a special case, the discontinuous Galerkin method.

To illustrate the efficacy of the proposed approach for the solution of conservation laws, we present examples of solving the three-dimensional Maxwell's equations for time-domain scattering from very complex bodies and the two-dimensional compressible Euler equations for two-fluid problems.

While the use of  $n$ -simplices ensures significant geometric flexibility, the nodal based approximation has a further advantage in that it enables the construction of very efficient preconditioners by allowing for the use of finite element-based preconditioning. We shall illustrate this latter point by examples of solving the two-dimensional incompressible Navier-Stokes equations emphasizing the need to consider conditioning techniques to remove the  $h$  as well as the  $p$  dependency. We shall offer some suggestions on how to accomplish this within the proposed framework.

The work discussed is done in collaboration with Dr. Tim Warburton (Brown University) and Dr. Luca Pavarino (University of Milan, Italy).

*Speaker's web page:* <http://www.cfm.brown.edu/people/jansh/>

*Research web page:* <http://www.cfm.brown.edu/home.html>

*Institution web page:* <http://www.brown.edu/>

March 24, 2000

# Charon Parallelization Toolkit

**Rob Van Der Wijngaart**

MRJ Technology Solutions, NASA  
Ames Research Center

Email: [wijngaar@nas.nasa.gov](mailto:wijngaar@nas.nasa.gov)

## Abstract:

**M**essage passing is among the most efficient methods for parallelization of scientific programs. Other strategies are available, but these often fail to deliver the flexibility, scalability, and high performance required for large-scale computations. This is especially true for applications developed for distributed-memory machines. Unfortunately, writing message-passing programs is a slow, cumbersome, and error-prone process.

Charon is an application-level library (C and Fortran) conceived to help ease the burden of the message-passing implementation of a class of 'difficult' parallel programs. This class is characterized by implicit numerical algorithms on structured grids that are not—or not easily—amenable to data-parallel solution techniques.

Charon supports piecemeal conversion of a serial design or legacy code to a fully parallel message-passing code by decoupling data distribution and parallel execution. This is achieved by tools that mimic serial program execution on distributed data. Through a sequence of tuning steps, again using tools provided by Charon, performance is gradually improved.

In this seminar the basics of the library will be discussed, and some examples of its use and performance will be presented.

*Institution web page:* <http://www.arc.nasa.gov/>



March 23, 2000

# Programming Tools for Very Large Dataset Subsetting and Aggregation

**Joel Saltz**

University of Maryland

Email: saltz@cs.umd.edu

## *Abstract:*

Increasingly powerful computers, clusters, and multiprocessor machines have enabled computational scientists and engineers to model physical phenomena in great detail. As a result, overwhelming amounts of data are being generated by scientific and engineering simulations. In addition, large amounts of data are being generated by sensors of various sorts, attached to devices such as satellites and microscopes. The exploration and analysis of the resulting large datasets plays an increasingly important part in many domains of scientific research. In this presentation we describe the design and development of software systems designed to address the need to subset, explore, analyze, process, and visualize very large datasets.

The first software system, the Active Data Repository (ADR), targets large disk-based datasets in processing environments with multiple processors and multiple disks. ADR is a C++ class library designed to support the development of servers that support multidimensional range queries with user defined aggregation and filtering functions.

The second, closely related, software system is a set of middleware infrastructure, called DataCutter, that supports the processing of scientific data collections stored in archival storage systems across a wide-area network. DataCutter also provides support for subsetting of datasets through multidimensional range queries along with support for invoking a sequence of user-defined filtering and aggregation functions. Processing, network, and data copying overheads are minimized by the ability to place filtering and aggregation functions on different platforms.

We will also discuss a compiler infrastructure that allow users to target ADR and DataCutter by specifying data distributions using XML and by writing user defined aggregation functions in a high-level language.

*Institution web page: <http://www.umd.edu/>*

March 16, 2000

*Abstract:*

# Data Mining: An Overview

**N. Radhakrishnan and  
Raju Namburu**

U.S. Army Research Laboratory (ARL)

**T**his talk will cover, in a light-hearted way, the emergence of data mining (a subset of Knowledge Discovery) in academia, industry, and the government. Questions like, what is data mining?, why is it needed?, how is it done?, etc., will be outlined. Several successful applications of data mining in industry and the government will be covered with special emphasis on scientific data mining. The talk will also highlight data mining advantages and pitfalls and the need for responsible management of data.

N. Radhakrishnan is the Director of the Corporate Information and Computing Directorate, and R. Namburu is the DOD thrust area lead for computational mechanics.

*Research web page:* <http://www.arl.mil/cicd/>

*Speaker's web page:* <http://www.arl.mil/cicd/radha.html>

*Institution web page:* <http://www.arl.mil/>



March 14, 2000

*Abstract:*

# Programming Environments for Multi- Application Simulation

**Eric DeSturler**

University of Illinois, Urbana-  
Champaign

Email: [sturler@uiuc.edu](mailto:sturler@uiuc.edu)

**T**he Center for Simulation of Advanced Rockets (CSAR) aims to develop simulation programs for the detailed, whole-system simulation of solid propellant rockets under normal and abnormal operating conditions. For the simulation of solid propellant rockets we need to solve highly complicated multi-physics problems, involving among others, computational fluid dynamics with particles and chemical reactions, turbulence, solid dynamics and combustion, and crack propagation.

Rather than incorporate the equations of the constituent problems in a single system and combine the technology of currently separate programs in a single program, we aim to continue to use the existing programs for each constituent part. We will solve each application separately. The mutual influence of connected structures and/or coupled physical and chemical processes is accounted for explicitly by coupling the simulations of all relevant structures and processes at their "boundaries." This preserves the efforts in existing programs and it allows the developer to make or keep the optimal choices (model, spatial discretization, time integration, and solution procedures) for each individual simulated process. The complication is, of course, how to couple the simulated processes in a physically and mathematically correct way and in a computationally efficient way, and how to run the various programs in concord. We aim to develop a programming environment that will support combining with minimal changes and in a short time frame multiple grid-based applications for complex multi-physics simulations, especially involving highly dynamic adaptive simulations. Such a programming environment will be useful for a wide range of applications besides rocket simulation.

I will give an overview of the work at the CSAR, our current efforts, and future plans for the programming environment.

*Institution web page:* <http://www.uiuc.edu/>

March 13, 2000

*Abstract:*

# Source Code Engineering Using SNiFF+ for Large-Scale Software Development

**Gil Chita**

TakeFive Software

Email: [gil@takefive.com](mailto:gil@takefive.com)

**S**NiFF+ is an interactive source code analysis tool. Gil Chita, a Field Application Engineer for Wind River Systems, Inc., will show how to use SNiFF+ to understand the workings of existing, unfamiliar code. The presentation will have three parts: a 10-minute slide show introduction and overview, a 40-minute demonstration of SNiFF+, and a 10-minute question and answer session. The presentation is targeted at engineers and scientists that need to write new software modules, integrate the modules into existing systems, and port them to different operating systems.

*Institution web page:* <http://www.takefive.com/>



March 7, 2000

*Abstract:*

# Performance Instrumentation and Visualization with SvPablo

**Luiz De Rose**

IBM TJ Watson Research Center

Email: [laderose@us.ibm.com](mailto:laderose@us.ibm.com)

In this talk I will present the Source View Pablo, a language and architecture independent toolkit for performance instrumentation, visualization, and analysis of programs written in C, Fortran 77, Fortran 90, and HPF, on sequential and parallel systems. Besides capturing application data via software instrumentation, SvPablo exploits multi-platform hardware performance counters, via its interface with PAPI, in order to capture the interaction of software and hardware. During execution of the instrumented code, the SvPablo library captures data and computes performance metrics on the execution dynamics of each instrumented construct on each processor. Because only statistics, rather than detailed event traces, are maintained, SvPablo can handle measurements of programs that execute for hours or days on hundreds of processors. In addition, SvPablo can be easily extended to new contexts with minimal changes to the software infrastructure.

*Research web page:* [www.research.ibm.com/actc/](http://www.research.ibm.com/actc/)

March 3, 2000

*Abstract:*

# Modeling the Performance of Sn Transport on Parallel Computers

**Harvey Wasserman**

Los Alamos National Laboratory

Email: [hjw@lanl.gov](mailto:hjw@lanl.gov)

**P**article transport via deterministic approaches such as the Discrete Ordinates Sn consumes enormous amounts of time during simulations of interest to ASCI. Moreover, the proportion of time spent in this application is expected to grow as ASCI begins to approach its canonical goal of order billion-cell problems. This talk will describe attempts to create a predictive capability for performance of a Cartesian-grid Sn transport application. The performance model attempts to parameterize characteristics of the machine (such as message latency and bandwidth and a characteristic floating-point computation rate) and of the application (such as number of "sweeps" and blocking factors for non-parallel computation). The talk will also describe attempts to generalize the model from a uniform network topology to lower-dimensional networks such as those used in clusters of SMPs.

*Speaker's web page:* <http://www.c3.lanl.gov/~hjw/hjw.html>

*Research web page:*

[http://www.c3.lanl.gov/cic19/teams/par\\_arch/par\\_arch.shtml](http://www.c3.lanl.gov/cic19/teams/par_arch/par_arch.shtml)

*Institution web page:* <http://www.lanl.gov/>



March 2, 2000

*Abstract:*

# Hybrid MPI-OpenMP Implementation of an Unstructured Grid Agglomeration Multigrid Solver

**Dimitri Mavriplis**

ICASE NASA Langley Research Center

Email: [dimitri@icase.edu](mailto:dimitri@icase.edu)

A two-level MPI-OpenMP parallel implementation of an existing unstructured multigrid Navier-Stokes solver is described. For both programming models, parallelism is achieved through domain partitioning applied to the various multigrid levels. In the MPI programming model, each process operates on a single partition and communication between the partitions is handled by explicit message-passing routines. In the OpenMP model, loops over the number of domains are parallelized using compiler directives, thus spawning one thread for each partition, and communication is handled by copying values from ghost vertices in individual partitions to their real images in adjacent partitions. The current implementation allows the solver to be executed either in a pure MPI mode, a pure OpenMP mode, or a two-level MPI-OpenMP mode. In the latter case, two hybrid communication strategies for combining MPI and OpenMP have been implemented. In the first approach, MPI messages are sent and received by individual OpenMP threads running in parallel. In the second approach, the thread-level messages are packed into larger messages which are then sent to remote MPI processes using the master thread, where they are subsequently unpacked.

Performance comparisons of MPI, OpenMP, and combined MPI/OpenMP using both communication strategies are given on the Cray-SV1, and the SGI Origin-2000. Benchmark results are demonstrated on all three current open ASCI machines using up to 1024 and 2048 processors.

The two-level parallel unstructured multigrid Navier-Stokes solver described herein will form the basis of a radiation diffusion solver, which is currently under development in two dimensions.

*Institution web page:* <http://www.larc.nasa.gov/>

February 28, 2000

*Abstract:*

# SCM by Perforce

**Ines Heinz**

Lawrence Livermore National  
Laboratory

Email: [heinz2@llnl.gov](mailto:heinz2@llnl.gov)

**P**erforce is a software configuration management system that was designed to provide a great deal of control while presenting a minimal footprint. The architecture of Perforce will be presented along with a brief comparison to ClearCase and CVS. Some strengths and weaknesses will be outlined. Finally, there will be a brief discussion of features and personal experiences.

*Institution web page:* <http://www.llnl.gov/>

February 24, 2000

# Fast Solvers in Scientific Computing

**Gabriel Wittum**

University of Heidelberg

Email: [wittum@iwr.uni-heidelberg.de](mailto:wittum@iwr.uni-heidelberg.de)

## *Abstract:*

**S**olvers play a crucial role in many simulations determining the complexity of the overall process. Thus, solving often limits the obtainable accuracy and fast solvers are the key to simulating new challenging problems.

In this lecture, numerical strategies for the simulation of large problems are discussed. The main directions are adaptivity, multigrid, and parallelism. The software system UG is presented which is based on these strategies. In several application cases the efficiency of the selected approach is shown.

*Research web page:* <http://www.iwr.uni-heidelberg.de/~techsim/>

*Institution web page:* [http://www.uni-heidelberg.de/index\\_e.html](http://www.uni-heidelberg.de/index_e.html)

February 23, 2000

*Abstract:*

# An Asymptotic Analysis of Multigroup Neutron Diffusion

**Gerald Hedstrom**

LLNL Retiree

Email: [ghedstrom@home.com](mailto:ghedstrom@home.com)

**I**t has been known for many years that for monoenergetic neutrons the diffusion equation is an asymptotic limit of the Boltzmann equation, and this information is used in the design of preconditioners for numerical methods. In this talk, I use the theory of positive operators to extend the results to the energy-dependent case.



February 11, 2000

*Abstract:*

# Compiler Technology and Scientific Computing

**David Padua**

University of Illinois

Email: [padua@uiuc.edu](mailto:padua@uiuc.edu)

Scientific computing has always been an important focus for compiler technology, to the point that the first commercial compiler was a Fortran compiler and many of the early optimization techniques were developed to improve the performance loops containing array accesses. Today, powerful compilers are an important part of any development environment for numerical codes. These compilers usually generate very good code that can take advantage of the capabilities of the most powerful processors.

Although much has been achieved since the first Fortran compiler, there are still many interesting and important issues in compiler technology that need attention. For example, techniques to analyze accesses to irregular data structures are not well understood although they are of great importance for conventional optimizations as well as detection of instruction level parallelism. Effective algorithms to handle explicitly parallel programs are also important because parallel programming is likely to become more frequent as SMPs grow in popularity. Today's techniques were developed for sequential programs and, therefore, do not always work correctly in a parallel context. For example, some compilers eliminate as dead code spin-lock loops, which thereby produces incorrect code and creates difficulties for the unsuspecting programmer.

Also, techniques to detect coarse grain parallelism are quite important and more needs to be done in this area, although the problem has proven much more difficult than originally expected. Other two issues are the systematization of the optimization process, which traditionally has been based on ad-hoc procedures that are not well understood, and the development of compiler algorithms that take advantage of information available in very high level programs to perform advanced optimizations.

In this talk, I will present a brief description of ongoing work at Illinois on these five topics.

*Speaker's web page:* <http://www.cs.uiuc.edu/contacts/faculty/padua.html>

*Research web page:* <http://polaris.cs.uiuc.edu/polaris/polaris.html>

*Institution web page:* <http://www.uiuc.edu/>

January 28, 2000

*Abstract:*

# Incremental Nearest Neighbor Finding

**Hanan Samet**

University of Maryland, College Park

Email: [hjs@umiacs.umd.edu](mailto:hjs@umiacs.umd.edu)

An algorithm for finding nearest neighbors incrementally is presented. It finds use in ranking spatial objects according to increasing distance from a query object in a graphical database. The algorithm makes use of a hierarchical spatial data structure. The algorithm is general in the sense that it is applicable to all hierarchical spatial data structures, all spatial data types, and to data of arbitrary dimensionality. The intended application area is a database environment, where the spatial data structure serves as an index. The algorithm is incremental in the sense that objects are reported one by one, so that a query processor can use the algorithm in a pipelined fashion for complex queries involving proximity. It is well suited for  $k$  nearest neighbor queries, and has the property that  $k$  needs not be fixed in advance. The algorithm has been used as the basis of a browser for graphical objects in a relational database. It also serves as the heart of the VASCO spatial data structure applet which will be presented in the talk (found at <http://www.cs.umd.edu/~hjs/quadtree/index.html>). Other applications of the algorithm include a distance semi-join which enables the computation of a discrete Voronoi diagram in an incremental manner. This is the result of joint work with Gisli R. Hjaltason.

*Speaker's web page:* <http://www.cs.umd.edu/~hjs/>

*Institution web page:* <http://www.umcp.umd.edu/>



January 27, 2000

*Abstract:*

# Chemometric Analysis of ICP-AES Spectroscopic Data Collected on Acidic Aqueous Samples Mimicking Nist Plain Carbon Steels

**John Quagliano**

Los Alamos National Laboratory

Email: JOHNQ-MAN@prodigy.net

There is considerable interest in using Inductively Coupled Plasma Atomic Emission Spectroscopy to characterize the minor elemental components of stainless steel. Iron interferes spectrally with many of the emission lines of the other analytes. Tedious and laborious wet chemical separations (to remove Fe) were previously required to obtain quantitative determinations. The present work demonstrates the application of a multivariate statistical analysis for the determination of six transition metals in a high concentration aqueous iron matrix. Sample preparation involved only a simple nitric acid dissolution. A 60 sample calibration set with a solvent background subtraction yielded a Partial Least Squares (PLS) model with 9 principal components after a leave-one-out cross validation. After independent validation, all analytes but vanadium were predicted with errors of no worse than 11%. Based on this preliminary investigation, we have learned new implementation strategies that should enable us to improve the quality of our results.

*Institution web page:* <http://www.lanl.gov/>

January 27, 2000

*Abstract:*

# Demonstration of Improved Analytical Chemical Operations Using the Analyst-Interactive Program FRAD (Fast Reduction of Analytical Data) Running Under Windows 95/98/NT an Application to ICP Mass Spectrometry

**John Quagliano**

Los Alamos National Laboratory

Email: [JOHNQ-MAN@prodigy.net](mailto:JOHNQ-MAN@prodigy.net)

**M**odern analytical instrumentation is capable of producing large data sets in a short amount of time due in part to multi-channel detection and computer automation. However, it is still the analyst's responsibility to interpret and organize the raw data into a useful and meaningful format and also to assure that numerous Quality Control (QC) criteria are satisfied. It has been our experience that instrument vendor and third party software programs are inadequate for these data reduction needs. Consequently, our chemists spend an inordinate amount of time performing spreadsheet manipulations to complete the data reduction. Therefore, there is much less time spent in the lab developing improved measurement techniques and the analytical services team produces a slower turnaround of results to customers.

To solve this problem, we have developed a stand-alone executable software program that formats raw data, performs QC checks, notifies the analyst of errors, prepares data for LIMS, and generates customer deliverable reports. The program FRAD is mouse/menu driven with the familiar Windows Graphical User Interfaces (GUIs) and is suitable for both technicians and scientists or engineers. As an example, this presentation will demonstrate how a 6-isotope ICP-MS raw data file that previously required 2 hours to process should now take less than half an hour to process without any loss of thoroughness or quality.

Our ultimate goal is to have this conveniently formatted set of data results serve as input for the commonly used and already available chemometric software packages. In this way, the combined FRAD/Chemometrics software can (a) make the basic chemometric tools more attractive to analytical service personnel and (b) assist R&D personnel in assessing the quality of their research data before treatment by advanced chemometric tools.

*Institution web page:* <http://www.lanl.gov/>



January 24, 2000

*Abstract:*

# Parallel Program Analysis Framework for the DOE ACTS Toolkit

**Allen Malony**

University of Oregon

Email: [malony@cs.uoregon.edu](mailto:malony@cs.uoregon.edu)

The diversity of high-performance computing platforms, programming systems, and applications and libraries that make up the computing environments at the DOE laboratories requires a rich set of tools. It is difficult, however, to build tools that are both portable across the set hardware and software systems, and able to be integrated with other computing and program components. This talk describes our work in developing program analysis technology for static and dynamic analysis of parallel programs. The work is funded by the DOE 2000 program and targets the analysis requirements of the ACTS toolkit.

Three project activities are covered: TAU, PDT, and the TAU distributed monitor. The TAU (Tuning Analysis Utilities) activity is developing a robust performance analysis framework that has been demonstrated on all ASCI platform types and across all ACTS software layers. The PDT (Program Database Toolkit) activity is developing a code analysis system that can be used to build source-level tools. Finally, the TAU distributed monitoring system is enabling runtime access to TAU performance data.

This work will be described in detail and future activities will be discussed.

*Research web page:* <http://www.cs.uoregon.edu/research/paracomp/proj/tau/>

*Institution web page:* <http://www.uoregon.edu/>

January 19, 2000

# Support for Dynamic Thread Management in Clustered Systems

**Steve Chapin**

Syracuse University

Email: [chapin@ecs.syr.edu](mailto:chapin@ecs.syr.edu)

## *Abstract:*

**T**he Heterogeneous Operating System Software (HOSS) project is building operating system software for high-performance, heterogeneous clusters. Although these clustered environments are capable of sustaining computation rates that rival or surpass conventional supercomputers, the performance delivered to scientific applications is only a fraction of this maximum. The intent of this research is to develop mechanisms for enabling high-level software such as applications, message-passing libraries, and runtime systems to better manage resources through the operating system kernel. The resulting operating system will provide management functions allowing middleware to predict the performance ramifications of policy choices, and to allow runtime systems to make quality-of-service guarantees to application programs.

*Speaker's web page:* <http://www.cs.virginia.edu/~chapin/>

*Institution web page:* <http://www.syr.edu/>



January 14, 2000

*Abstract:*

# An Architecture for Management and High- Performance Access of Data

**Alok Choudhary**

Northwestern University

Email: [choudhar@ece.nwu.edu](mailto:choudhar@ece.nwu.edu)

**L**arge-scale scientific applications (e.g., ASCI applications) consume and produce tremendous amounts of data for analysis, visualization, and other purposes. It can be an overwhelming task just to manage this data for a large number of experiments. This problem is further exacerbated when performance is critical (in particular I/O), sharing data among multiple applications is required, and hierarchical and distributed storage systems are used. The number of system and performance-related parameters can be large, thereby making performance optimizations for I/O a daunting task where a user has to understand what optimizations are suitable under what conditions and how to achieve them.

In this talk we will present an architecture and initial implementation for a system for managing scientific data with a particular emphasis on “performance management” and “automatic performance optimizations” for data access. A DBMS is used to manage performance metadata which can store and use data access patterns and storage patterns, but data access is done using a runtime system outside the DBMS. In particular we will present (1) the architecture and design philosophy, (2) the runtime system and automatic access optimizations using the database model, (3) initial performance results, (4) initial results for HPSS access patterns, and (5) user interface for management using runtime system. We will also summarize other work as part of our project.

*Speaker's web page:* <http://web.ece.nwu.edu/~choudhar/>

*Institution web page:* <http://www.nwu.edu/>

January 13, 2000

# The PARASOL Project: Overview and Domain Decomposition Solvers

**Petter Bjorstad**

University of Bergen, Norway

Email: [petter@ii.uib.no](mailto:petter@ii.uib.no)

## *Abstract:*

**P**art I: An overview of the PARASOL sparse linear solvers. The Parasol project, an R&D project supported by the European Commission from 1996 through the summer of 1999, will be described, with a special focus on the sparse, linear solver developments.

Part II: Experiences with Domain Decomposition Algorithms in the PARASOL project. This talk will describe the effort to create a framework for Domain Decomposition algorithms (or preconditioners) within the PARASOL project, a large R&D project on industrial strength, parallel, and sparse linear solvers.

The main emphasis will be on issues related to a robust implementation of the Neumann–Neumann algorithm for realistic problems resulting from finite element discretization. We will discuss the construction of appropriate coarse spaces and several issues related to hybrid sparse algorithms, that is, the combination of direct and iterative sparse methods.

*Speaker's web page:* <http://www.ii.uib.no/~petter/>

*Research web page:* <http://www.parallab.uib.no/>



January 11, 2000

*Abstract:*

# Very High Resolution Simulation of Compressible Turbulence on the IBM-SP System

**Art Mirin**

Lawrence Livermore National Laboratory

**Winner of 1999 Gordon Bell Award for Performance**

A. Mirin, R. Cohen, B. Curtis, W. Dannevik, A. Dimits, M. Duchaineau, D. Eliason and D. Schikore, LLNL, S. Anderson, D. Porter and P. Woodward, University of Minnesota, L. Shieh and S. White, IBM

Understanding turbulence and mix in compressible flows is of fundamental importance to real-world applications such as chemical combustion and supernova evolution. The ability to run in three dimensions and at very high resolution is required for the simulation to accurately represent the interaction of the various length scales, and consequently, the reactivity of the intermixing species. Toward this end, we have carried out a very high resolution (over 8 billion zones) 3-D simulation of the Richtmyer-Meshkov instability and turbulent mixing on the IBM Sustained Stewardship TeraOp (SST) system, developed under the auspices of the Department of Energy (DOE) Accelerated Strategic Computing Initiative (ASCI) and located at Lawrence Livermore National Laboratory.

We have also undertaken an even higher resolution proof-of-principle calculation (over 24 billion zones) on 5832 processors of the IBM system, which executed for over an hour at a sustained rate of 1.05 Tflop/s, as well as a short calculation with a modified algorithm that achieved a sustained rate of 1.18 Tflop/s. The full production scientific simulation, using a further modified algorithm, ran for 27,000 timesteps in slightly over a week of wall time using 3840 processors of the IBM system, clocking a sustained throughput of roughly 0.6 teraflop per second (32-bit arithmetic). Nearly 300,000 graphics files comprising over three terabytes of data were produced and post-processed.

The capability of running in 3-D at high resolution enabled us to get a more accurate and detailed picture of the fluid-flow structure—in particular, to simulate the development of fine-scale structures from the interactions of long- and short-wavelength phenomena, to elucidate differences between two-dimensional and three-dimensional turbulence, to explore a conjecture regarding the transition from unstable flow to fully developed turbulence with increasing Reynolds number, and to ascertain convergence of the computed solution with respect to mesh resolution.

*Speaker's web page:* <http://www.llnl.gov/CASC/people/mirin/>

*Project web page:* <http://www.llnl.gov/CASC/asciturb/>

January 6, 2000

*Abstract:*

# Parallel Implementation of Multigrid Methods on Unstructured Grids using Adaptive Finite Element

**Linda Stals**

Old Dominion University

Email: [stals@cs.odu.edu](mailto:stals@cs.odu.edu)

**S**olving systems of equations using the multigrid method on a parallel machine presents a challenge due to the need to handle both the intra-grid and inter-grid data dependencies. Our approach to tackling such complicated data dependencies is to implement a very flexible node-edge data structure using an object-oriented framework. This setup has allowed us to write the code in a manner which more closely mimics our conceptual view of the problem and thus we are able to refine the grids (in 2D and 3D), build the system of equations, balance the load, and solve the system of equations all in parallel.

During the presentation I plan to present examples from the solution of the radiation transport equations (ODU and ICASE). I shall also briefly mention some other applications, including the Plasma Ion Immersion Process (Universitat Augsburg) and flow in heterogeneous media (University of Bath).

*Speaker's web page:* <http://www.icas.edu/~stals/>

*Research web page:* <http://www.icas.edu/~keyes/asci/>

*Institution web page:* <http://web.odu.edu/>

December 14, 1999

# Scheduling and Processor Allocation for Matrix Chain Products on Parallel Systems

**Jong Kim**

Pohang University of Science and Technology

Email: [jkim@postech.ac.kr](mailto:jkim@postech.ac.kr)

*Abstract:*

**M**atrix computations are fundamentals of scientific computing and commonly arising computationally intensive parts of many numerical and non-numerical applications. In this talk, I present the computation of matrix chain products on parallel systems.

The problem we are considering is finding an optimal product schedule for evaluating a chain of matrix products on a parallel computer (the matrix chain scheduling problem, MCSP). Since the approach of parallelizing each matrix product after finding an optimal product sequence for single processor systems does not always guarantee the minimum evaluation time on parallel systems, we introduce a new processor scheduling and processor allocation algorithms for the MCSP which reduces the evaluation time of a chain of matrix products on a parallel computer. We show that the proposed scheduling algorithm significantly decreases the time required to evaluate a chain of matrix products in parallel systems. The scheduling approach may also improve the performance of many other scientific computations.

*Speaker's web page:* <http://www.postech.ac.kr/~jkim/>

*Institution web page:* <http://www.postech.ac.kr/e/>

December 13, 1999

# The Urgent Need for Configuration Management

**Michael Donaldson**

Merant

Email:

Michael.Donaldson@merant.com

## *Abstract:*

**C**onfiguration management—the disciplined approach to managing the evolution of software development and maintenance practices and their products—has matured to the point where organizations now cannot thrive without configuration management (CM) capabilities. Industry has learned many lessons from its software crisis: namely, that CM is a mission-critical, foundational practice. Organizations have seen the value of a full-featured process-oriented CM solution. They include:

- Improved release cycle times
- Reduction in bugs
- Automatic quality control
- Minimal change complexity
- Variant/concurrent release management
- Visibility into all work status
- Better forecasting of release dates
- Increased preparedness for audits
- Repeatability of processes and product creation
- Improved responsiveness to customers
- Faster standards certifications

PVCS Dimensions is a full-featured, process-oriented CM solution that provides systemic quality automation by operationally supporting: version control, parallel and distributed development, build and release management, workspace and repository management, change and issue tracking, status reporting, and audit control.

This presentation discusses why CM is urgently needed today, the role CM plays in process improvement goals, and how PVCS Dimensions maximizes your investment in software development.

*Institution web page: <http://www.merant.com/>*

December 3, 1999

*Abstract:*

# Smarter Memory Controllers: Improving Memory System Performance from the Bottom Up

**Sally McKee**

University of Utah

Email: [sam@cs.utah.edu](mailto:sam@cs.utah.edu)

**M**icroprocessor speed is increasing much faster than memory system speed: both speed-growth curves are exponential, but they represent diverging exponentials. The traditional approach to attacking the memory system bottleneck has been to build deeper and more complex cache hierarchies. Although caching may work well for parts of programs that exhibit high locality, many important commercial and scientific workloads lack the locality of reference that makes caching effective. A 1996 study by Sites and Perl on a commercial database workload shows that memory bus and DRAM latencies cause an 8X slowdown from peak performance to actual performance. Another 1996 study by Burger, Goodman, and Kagi finds that when compared to an optimal cache the efficiency of current caching techniques is generally less than 20% and that cache sizes are up to 2000 times larger. The evidence is clear: no matter how hard we push it, traditional caching cannot bridge the growing processor-memory performance gap.

This talk presents research that attacks the memory problem at a different level—the memory controller. We describe the Stream Memory Controller (SMC) system built at the University of Virginia, and the Impulse Adaptable Memory System being built at the University of Utah. The SMC dynamically reorders accesses to stream elements to avoid bank conflicts and bus turn-around delays and to exploit locality of reference within the page buffers of the DRAMs. Impulse virtualizes unused physical addresses and dynamically remaps data to improve cache and bus utilization. For instance, Impulse can gather sparse data into dense cache lines that have high locality and leave a smaller cache footprint. Gathering strided data improves performance on regular applications, and gathering data through an indirection vector improves performance on irregular applications (e.g., Impulse improves the NAS conjugate gradient benchmark's performance by 67%). By prefetching data within the memory controller, Impulse uses bus bandwidth more efficiently, transferring only the needed data when it is requested.

Both these systems require that the compiler or application writer supply the access pattern information that the memory controller exploits. The SMC optimizes low-level memory performance for regular applications. In contrast, Impulse optimizes performance within the cache hierarchy and the memory back end for both regular and irregular computations.

*Speaker's web page:* <http://www.cs.utah.edu/~sam/>

*Institution web page:* <http://www.utah.edu/>

December 2, 1999

# Scalable Network Attached Secure Storage Systems

**Garth Gibson**

Carnegie Mellon University

Email: [garth@cs.cmu.edu](mailto:garth@cs.cmu.edu)

## *Abstract:*

While it is possible to construct off-the-shelf, widely distributed, and massively parallel storage systems with inherent high bandwidth, achieving low-latency file access remains a significant challenge. We are designing, implementing, and evaluating scalable, distributed, and parallel storage architectures, interfaces, and protocols to reduce access latency comprehensively. Our goal is to define the evolutionary path and revolutionary changes that will enable commodity storage components to be the building blocks of high-bandwidth, low-latency, secure scalable storage systems.

Our current definition for a NASD device includes all storage systems that exhibit direct client-drive data transfer in a networked environment and asynchronous oversight by the high-level filesystem.

*Institutional web page:* <http://www.cmu.edu/>



November 23, 1999

*Abstract:*

# Quality Issues in Image Compression: Perceptual Measures and Processing for Performance Improvement

**V. Ralph Algazi**

University of California, Davis

Email: [algazi@ece.ucdavis.edu](mailto:algazi@ece.ucdavis.edu)

Two broad issues that determine the quality of compressed images are the perceptual importance of details or noise in the original image, and the artifacts introduced by the compressor. In this presentation, an overview and discussion of subjective and objective quality measures in the compression of achromatic still images are presented. We first consider subjective quality scales and measures based on properties of human perception, and then present and discuss some objective measures and their correspondence to the subjective assessment of quality. An interesting alternative to the compression of the original image is to process the image prior to compression, so as to reduce unimportant details while improving the compressibility of the image data. Some techniques for improving the bit-rate versus distortion trade-off by image processing are presented, as well as the quality measures that allow to quantify the performance improvement possible.

*Speaker's web page:* <http://info.cipic.ucdavis.edu/~algazi/index.html>

*Institution web page:* <http://www.ucdavis.edu/>

November 22, 1999

# Generic Software Components for Scientific Computing

**Jeremy Siek**

University of Notre Dame

Email: [jsiek@lsc.nd.edu](mailto:jsiek@lsc.nd.edu)

## *Abstract:*

**D**eveloping high-quality software is an exceedingly difficult task. Although software engineers continue to develop new methodologies that promise to enable the development of high-quality software, many of these methodologies do not even live up to their own promises, much less to the more stringent requirements of scientific applications. All is not lost, however. Careful study of various software development paradigms in the context of high-quality scientific software clearly points to several new approaches that demonstrate significant promise both in theory and in practice. We validate our approach with the recently developed Matrix Template Library (MTL), a comprehensive collection of generic software components for high-performance scientific computing. Four aspects of the design and development of MTL will be presented. First, we discuss the generic programming methodology used to create the MTL framework. Second, we describe the framework of the abstract MTL interface. Third, we present the concrete implementation of MTL. Finally, we give performance results demonstrating that the generic components MTL are able to achieve performance equal to that of vendor-tuned (and automatically tuned) numerical libraries.

*Speaker's web page:* <http://www.lsc.nd.edu/~jsiek/>

*Research web page:* <http://www.lsc.nd.edu/research/mtl/>

*Institution web page:* <http://www.nd.edu/>



November 22, 1999

# The Generic Graph Component Library

**Andrew Lumsdaine**

University of Notre Dame

Email: [lums@lsc.nd.edu](mailto:lums@lsc.nd.edu)

## Abstract:

We present the Generic Graph Component Library (GGCL), a generic programming framework for graph data structures and graph algorithms. Following the theme of the Standard Template Library (STL), the graph algorithms in GGCL do not depend on the particular data structures upon which they operate, so that a single algorithm can operate on arbitrary concrete representations of graphs. To attain this type of flexibility for graph data structures, which are more complicated than the containers in STL, we introduce several important concepts that form the generic interface between the algorithms and the data structures. We describe the principal abstractions comprising the GGCL, the algorithms and data structures that it provides, and we provide examples that demonstrate the use of GGCL to implement some common graph algorithms, as well as an implementation of the minimum degree sparse matrix ordering algorithm. Also discussed are particular implementation issues related to performance. Experimental results are presented to demonstrate that the novel lightweight implementation techniques and static polymorphism in GGCL results in code that is significantly more efficient than similar libraries written using the object-oriented paradigm. In the case of sparse matrix ordering, our implementation of the minimum degree algorithm has performance indistinguishable from that of a widely used (and highly tuned) Fortran implementation.

*Speaker's web page:* <http://www.lsc.nd.edu/~lums/>

*Research web page:* <http://www.lsc.nd.edu/research/ggcl/>

*Institution web page:* <http://www.nd.edu/>

November 12, 1999

# Uniform Data Access Using GXD

**Peter Vanderbilt**

MRJ Technology Solutions

Email: [pv@nas.nasa.gov](mailto:pv@nas.nasa.gov)

## *Abstract:*

**T**his talk will give an overview of GXD, a framework facilitating publication and use of data from diverse data sources. GXD defines an object-oriented data model designed to represent a wide range of things including data, its metadata, resources, and query results. GXD also defines a data transport language, a dialect of XML, for representing instances of the data model. GXD also provides a software library, prototyped in Java, that includes support for generating GXD-encoded entities and for interpreting these entities (potentially from many data sources) to create an illusion of a globally interconnected data space, one that is independent of data source location and implementation.

*Institution web page:* <http://www.mrj.com/>



November 2, 1999

# Dyninst: Status and Future Evolution

**Jeff Hollingsworth**

University of Maryland, College Park

Email: [hollings@cs.umd.edu](mailto:hollings@cs.umd.edu)

## *Abstract:*

I describe a post compiler program manipulation called dyninstAPI which provides a C++ class library for machine independent program instrumentation during application execution. After a brief introduction to the topic, the talk will concentrate on presenting the new and recently proposed features in the API including support for type information, program modification, and use as an off-line binary editor. In addition, I will describe a new TCL-based command line tool we have developed to speed the creation of simple tools that use the API. I will also discuss the status of the DPCL interface being developed by IBM.

*Speaker's web page:* <http://www.cs.umd.edu/~hollings/>

*Research web page:* <http://www.cs.umd.edu/projects/dyninstAPI/>

*Institution web page:* <http://www.umd.edu/>

October 20, 1999

*Abstract:*

# KULL System Architecture

**Pat Miller**

Scientific Computing Applications  
Division

Lawrence Livermore National  
Laboratory

**T**he KULL project is a joint A-Division/X-Division effort to build a 3-D modeling code for simulating ICF capsules. The code is built on an object-oriented framework. Individual physics packages are written either in C++ or in Fortran. We make clever use of templization and inheritance to promote code reuse and to maintain a very abstract view of the mesh geometries used by the physics packages. The entire code is “steered” from a Python interpreter with embedded physics objects. The presentation will include some architectural details of the project, including our use of parallelism, C++ features, and steering.

*Institution web page:* <http://www.llnl.gov/>



October 18, 1999

# Diamond Eye: A System for Mining Large Image Collections

**Michael Burl**

Jet Propulsion Laboratory

Email: [burl@aig.jpl.nasa.gov](mailto:burl@aig.jpl.nasa.gov)

## *Abstract:*

**D**iamond Eye is a distributed software architecture that enables users (scientists) to analyze large image collections by interacting with one or more custom data mining servers via a Java applet interface. Each server is coupled with an object-oriented database and a computational engine such as a network of high-performance workstations. The database provides persistent storage and supports querying of the “mined” information. The computational engine provides parallel execution of expensive image processing, object recognition, and query-by-content operations.

Key benefits of the Diamond Eye architecture are: (1) the design promotes trial evaluation of advanced data mining and machine learning techniques by potential new users (all that is required is to point a web browser to the appropriate URL); (2) software infrastructure that is common across a range of science mining applications is factored out and reused; and (3) the system facilitates closer collaborations between algorithm developers and domain experts.

*Speaker's web page:*

<http://www-aig.jpl.nasa.gov/public/mls/home/burl/home.html>

*Institution web page:* <http://www.jpl.nasa.gov/>

October 15, 1999

*Abstract:*

# Development of a Three- Dimensional Relativistic PIC Code for Studying the Production of Useful Electron Bunches

**Roy Hemker**

University of California, Los Angeles

Email: [hemker@physics.ucla.edu](mailto:hemker@physics.ucla.edu)

The advances in computational speed make it now possible to do full 3D PIC simulations of laser plasma and beam plasma interactions, but at the same time the increased complexity of these problems makes it necessary to apply modern approaches like object-oriented programming to the development of simulation codes. We report here on our progress in developing an object-oriented parallel 3D PIC code using Fortran 90. In its current state the code contains algorithms for 1D, 2D, and 3D simulations in cartesian coordinates and for 2D cylindrically symmetric geometry. For all of these algorithms the code allows for a moving simulation window and arbitrary domain decomposition for any number of dimensions. Recent 2D and 3D simulation results on the propagation of intense laser and electron beams through plasmas will be presented.

*Institution web page:* <http://www.ucla.edu/>



October 7, 1999

*Abstract:*

# Semiclassical Quantum- Mechanical Modelling using Parallel PIC Methods

**Dean Dauger**

University of California, Los Angeles

Email: [dauger@physics.ucla.edu](mailto:dauger@physics.ucla.edu)

**W**e have developed a multiparticle quantum-mechanical dynamic modeler based on a semiclassical approximation of Feynman path integrals. This approximation transforms the quantum problem into a problem of classical trajectories, thus enabling us to implement the simulation on parallel computers using mature Particle-In-Cell methods developed for plasma codes. With modern hardware, this should allow us to simulate multiple (100-1000) interacting quantum particles, treating the Hamiltonians semiclassically. Techniques that enable the numerical reliability of the simulation were developed to handle the grid-point nature of PIC methods and boundary conditions for the semiclassical case. One-dimensional single and multiparticle examples will be shown, with analyses of their energetic structure, and a technique for analyzing fermionic cases will be introduced.

*Institution web page:* <http://www.ucla.edu/>

October 5, 1999

# Nonlinear Solver Challenges in Computational Fluid Dynamics

**David Young**

The Boeing Company

Email: [david.p.young@boeing.com](mailto:david.p.young@boeing.com)

## *Abstract:*

**M**any important phenomena in fluid dynamics are highly nonlinear. Some examples include shock waves and boundary layer separation. There are well-known instances of non-uniqueness and ill-posedness even for inviscid transonic flow. These difficulties are not ameliorated by solving the Navier–Stokes equations. Design optimization problems inherit all these challenges and add ill-posedness related to constraints. In this talk, we will discuss some of these difficulties and some computational strategies for dealing with them. These strategies include nonlinear elimination, nonlinear Schwartz methods, and the use of multi-point design. A rather complete description of our design optimization methodology will be given with a discussion of the implications of various methodological choices for global convergence to the solution of the problem.

*Institution web page:* <http://www.boeing.com/>



October 4, 1999

*Abstract:*

# Nonlinear Solvers and Preconditioners in Groundwater Flow Simulations

**Timothy Kelley**

North Carolina State University

Email: [Tim\\_Kelley@ncsu.edu](mailto:Tim_Kelley@ncsu.edu)

**N**ewton–Krylov–Schwarz methods solve nonlinear equations by using Newton’s method with a Schwarz domain decomposition preconditioned Krylov method to approximate the Newton step. In this talk we will discuss the design and implementation of Newton–Krylov–Schwarz solvers in the context of the implicit temporal integration on an unstructured three-dimensional spatial mesh of Richards’ equation for flow in the unsaturated zone. The issues include nonsmooth nonlinearities, construction and efficient implementation of the coarse-mesh problem, and temporal integration. We conclude with a report of some preliminary numerical results.

*Speaker’s web page:* <http://www.math.ncsu.edu/~ctk/>

*Institution web page:* <http://www.ncsu.edu/>